

Planetary Robotic Exploration Driven by Science Hypotheses for Geologic Mapping

Alberto Candela¹, David Thompson², Eldar Noe Dobrea³, and David Wettergreen¹

Abstract—Planetary exploration involves frequent scientific reformulation and replanning. It is limited by communication constraints and to overcome this limitation, this paper formulates the process as a collaboration in which the human scientist and the robot work together to fill in gaps in knowledge to make discoveries. It introduces the science hypothesis map as the probabilistic structure in which scientists initially describe their abstract beliefs and hypotheses, and in which the state of this belief evolves as the robot makes raw measurements. It discusses how to incorporate path planning for maximizing scientific information gain, which is efficiently computed. As proof of concept, this paper describes a geologic exploration problem where a robot uses a spectrometer to infer the geologic composition of different regions in a mining district at Cuprite, Nevada. It shows that the science hypothesis map can infer geologic units with high accuracy, and that exploration using information gain-based path planning has better performance than exploration with conventional science-blind algorithms.

I. INTRODUCTION

Modern planetary robotic exploration is guided by scientists specifying the waypoints on a path that they believe will best address the investigation questions. The path is formed from expert knowledge of the site and expectations about where to gather mission-critical information. Scientists reinterpret their measurements with growing contextual knowledge of the environment, so real exploration is characterized by a frequent reformulation and replanning throughout the mission lifetime [1]. Replanning occurs on large strategic scales, bypassing or favoring geographic locales, as well as local tactical scales, lingering at an anomalous feature for additional measurements [2], [3]. However, many exploration scenarios occur under low bandwidth and high latency communications, leaving limited opportunities to revise the exploration plan.

This work describes an approach to overcome the communication bottleneck in robotic exploration, where the command to the remote explorer is based on an evolving model of what the initially scientist believes, rather than a single prescribed route, enabling the robot to take more adaptive and efficient actions based on real-time information, improving the rate and productivity of discovery.

This paper introduces the *science hypothesis map* as the spatial probabilistic structure in which scientists communi-




Fig. 1: Adaptive robotic exploration as a cycle of two processes: Bayesian inference and experimental design (path planning). Scientists can guide the process by defining hypotheses and constraints (in the form of resource models, goal waypoints, etc.).

cate their belief about the world and in which the belief state evolves as the robotic explorer collects information. This article develops the groundwork for science hypothesis maps, detailing how they can guide the robot to improve the rate and productivity of discovery.

In general, scientists define their exploration questions in terms of abstract concepts far removed from the raw sensor data available to a robot. It is usually unfeasible to encode this rich, abstract knowledge in a direct way that enables true robotic understanding. However, it is often possible to construct simpler hierarchical probabilistic models relating these representations to measurable data. Conventions for mission design show ways of quantifying these relationships. For example, NASA missions represent the relationship between abstract investigation objectives and raw measurements with a Science Traceability Matrix (STM). This is done through a tripartite division into *investigation objectives*, *physical properties*, and *instrument measurements*. Table I shows some examples of Science Traceability Matrices for geologic, oceanographic [4], and pollution mapping.

The science hypothesis map extends the idea of a Science Traceability Matrix to a probabilistic model with spatial extent, granting robustness under uncertainty. It also determines the information that collected measurements, directly interpretable by the robot, provide with respect to the fundamental investigation objective. The most beneficial measurements produce a large expected reduction in uncertainty. The previous notion serves as the foundation of *Bayesian experimental design* [5]. It can be decomposed into two processes that

¹A. Candela and D. Wettergreen are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA albertoc@andrew.cmu.edu, dsw@ri.cmu.edu

²D. Thompson is with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA David.R.Thompson@jpl.nasa.gov

³E. Noe Dobrea is with the Planetary Science Institute, 1700 East Fort Lowell, Suite 106, Tucson, AZ 85719, USA eldar@psi.edu

TABLE I: Examples of simplified Science Traceability Matrices.

Investigation objectives	Physical properties	Measurements
Map geologic formation processes in a region	Abundance of key mineral classes (Quartz, Kaolinite, Calcite, etc.)	Surface reflectance, 450-2500 nm, at 10 nm spectral resolution and >300 SNR
Map water quality in a river	Salinity of the water	Water conductivity in $\mu\text{S}/\text{cm}$
Map air quality in a city	Abundance of particles such as carbon monoxide, sulfur dioxide, etc.	Particle concentration in $\mu\text{g}/\text{m}^3$

form a cycle: *Bayesian inference* and *experimental design*. In the first one, the hypotheses' probabilities are updated when new measurements are collected. In the second one, the updated hypotheses dictate actions that minimize uncertainty, which translate to *path planning* in robotic exploration (Figure 1). With each communication event, the scientist can define the variables of interest and the probabilistic models relating them to instrument data. In real time, a robot can continually calculate navigation plans that optimize scientific discovery.

A few deployed robots can perform automatic science data analysis [2], [4], [6], [7], however they do not use science hypotheses as an integral part of their reasoning. Instead, they pursue static objectives that are fixed at the outset. These simplified tasks, such as mapping scalar fields (e.g., ocean temperature [4]), or detecting transient features (dust devils on Mars [8]), are defined long in advance. Some robots have used Bayesian networks for tasks such as mineral classification [9], or meteorite identification [10], but they also operate under predefined static objectives that ignore the evolution of the robot's overall knowledge of a scene throughout the mission. Researchers have only recently begun to investigate the complete Bayesian experimental design loop (Figure 1), and consequently build probabilistic models connecting high-level concepts with low-level observations [11], [12]. However, these approaches have not investigated the influence of the scientist's prior knowledge of a scene during exploration, which is critical in realistic scenarios. They have also not integrated nor compared state-of-the-art informative path planners with their systems.

The presented formulation has several specific benefits. First, science hypothesis maps improve operational efficiency. Current operations require meticulous direction of each robot. Science hypothesis maps communicate objectives simply and intuitively, and define the appropriate behaviors without low-level action planning. Second, they can improve science yield by forcing exploration to be driven by quantitative, formal hypotheses that are related mathematically to the measurements made by the robot. Too often, such hypotheses are left implicit or only heuristically tied to the activities selected. Robotic activity plans can be optimized, allowing the system to react to unanticipated events that occur while the robot is out of communication. For example, delays in navigation are common due to hazard avoidance [3] or unanticipated environmental influences. Robots that understand the hypotheses under study can recover from these effects by recomputing the measurement strategy. They can exploit resource surpluses in a similar fashion [3].




Fig. 2: Plate notation of the science hypothesis map. Each spatial region $r \in R$ has an independent hierarchical conditional distribution of investigation objectives $h \in H$, physical properties $y \in Y$, and measurements $z \in Z$.

This article first presents a general description of the science hypothesis map and the way it can be used to perform Bayesian inference. Next, it describes information gain-based path planning strategies that can be integrated with the science hypothesis map to optimize the rate and productivity of discovery. Finally, it demonstrates the application of these methods to a simulated geologic mapping problem of a well-studied site: the Cuprite Mining District, Nevada.

II. APPROACH

A. Science Hypothesis Map

The *science hypothesis map* has the following hierarchical components, visually represented in Figure 2:

- For simplicity, the map partitions into k independent, predefined spatial *regions*, labeled as $R : \{r_1, \dots, r_k\}$, (see left of Figure 2).
- The *investigation objectives* estimate abstract properties which are themselves unknowns, labeled as $H : \{h_1, \dots, h_l\}$. Each region can be explained by an investigation objective with probability $P_R(H)$.
- *Physical properties*, labeled as $Y : \{y_1, \dots, y_m\}$, have unique associations with the investigation objectives, given by $P_R(Y|H)$.
- The robot collects n sensor *measurements*, labeled as $Z : \{z_1, \dots, z_n\}$. These measurements could be noisy or dependent on observing conditions, so there is an indirect association between Z and Y given by $P_R(Z|Y)$.

For mapping convenience, the explored environment is partitioned into disjoint regions, each associated with its own independent conditional distributions relating investigation objectives, physical properties, and measurements (Figure 2).

B. Bayesian Inference

The next step is to use the previous probabilistic model to infer the corresponding investigation objective H for each region $r \in R$, given a measurement Z . That is, find a closed-form expression for $P_R(H|Z)$ in terms of the known conditional distributions $P_R(Y|H)$ and $P_R(Z|Y)$. After performing some standard Bayesian factorization operations, the expression is:

$$P_R(H|Z) \propto P_R(H) \sum_Y P_R(Y|H) P_R(Z|Y). \quad (1)$$

This is equivalent to performing Bayesian inference, with $P_R(H)$ being the prior probability or initial belief of the explaining investigation objective of a region, given by the scientist in advance, and $P_R(H|Z)$ the posterior probability or updated belief.

III. INFORMATIVE EXPLORATION

This paper focuses on the specific class of *path planning* where a mobile agent must explore an environment by planning an information-optimal path to reach a prescribed end-of-day goal location, optimally balancing navigation and resource costs against meaningful science measurements. During the mission, the robotic explorer collects a sequence of n measurements at spatial locations $X = (x_1, x_2, \dots, x_n)$, with the corresponding possible combinations of the n spectrum measurements as Z^n .

A. Information Gain Objective Function

The robot aims to optimize *information gain*, which is equivalent to *mutual information*. It is a classical objective function for information-driven action selection [13], [14]. Specifically, it is defined as the expected reduction in uncertainty after collecting new information, measured with Shannon entropy. The Shannon entropy of an independent region in R is:

$$I_R(H) = - \sum_H P_R(H) \log P_R(H). \quad (2)$$

The expected entropy of the posterior distribution given n measurements is:

$$I_R(H|Z^n) = - \sum_{Z^n} P_R(Z^n) \sum_H P_R(H|Z^n) \log P_R(H|Z^n). \quad (3)$$

From equations 2 and 3, the information gain for each region can be calculated as:

$$IG_R(H|Z^n) = I_R(H) - I_R(H|Z^n). \quad (4)$$

Information gain is additive across independent variables. Then, the objective function for the whole map can be represented as the sum of regions' information gains:

$$IG = \sum_R IG_R(H|Z^n). \quad (5)$$




Fig. 3: Information gain calculation for a uniform prior using three different methods: direct, Monte Carlo integration, and GLF fitting.

B. Information Gain Computation

A real time application typically faces the problem of calculating information gain efficiently. In this case, the sum over all possible sequences of n observations has exponentially many terms. While information gain can be estimated with Markov chain Monte Carlo approximations [15], here the distributions $P(Z^n)$ and $P(H|Z^n)$ are not known *a priori*. An alternative is to perform Monte Carlo integration [16], with the possibility of using simple sampling strategies. But another problem arises: the approximation accuracy deteriorates as n increases. However, given the problem formulation in this paper, there is a more robust and efficient approximation for large n . Information gain is a function that as long as Z and H are not independent, and the observations are conditionally independent given H , the following holds:

$$\lim_{n \rightarrow \infty} IG_R(H|Z^n) = I_R(H). \quad (6)$$

as shown by Haussler and Opper [17]. Since information gain (equations 3 and 4) is calculated by marginalizing Y (equation 1), these properties hold true.

This means that it is plausible to fit a curve to the first few data points (i.e. small number of measurements n), which are simpler to calculate or to approximate, specifically, a monotonically increasing curve that converges to a maximum. In this paper, a *generalized logistic function* (GLF) is used, also known as Richard's curve, which is a generalization of the sigmoid curve [18].

Figure 3 demonstrates an example of information gain calculation for a simple probabilistic process with a uniform prior. It uses nonlinear least squares fitting to find a suitable GLF approximation. It also shows a Monte Carlo approach, as well as the true information gain values. It can be easily observed that the accuracy of the Monte Carlo method deteriorates for large n , but not for the GLF method. In this example, only the first 6 Monte Carlo points are used for the GLF fit. All Monte Carlo approximations are calculated with 10,000 random points each by performing importance sampling. The proposal distribution is a uniform distribution over all possible combinations of Z^n , where points are sampled with the method by Smith and Tromble [19]. The goodness of fit statistics for $0 \leq n \leq 15$ are: $R^2 = 0.9999$

and $RMSE = 0.0092$, showing that the function can predict IG_R for many samples.

C. Information Gain-based Path Planning

This paper formulates the path planning objective as maximization of information gain subject to a fixed resource budget (time, energy, etc.) that the robot can travel before finally reaching a desired end-of-day goal. Implicitly, $x_1 = x_{start}$ and $x_n = x_{end}$. The optimization problem is:

$$\arg \max_X IG \quad \text{subject to} \quad Cost(X) \leq Budget. \quad (7)$$

Information gain is submodular. In this context, it means that sampling the same region repeatedly leads to diminishing expected returns. From the objective function in equation 5 and the submodularity of information gain, there is a resulting exploitation vs. exploration trade-off that translates to taking many samples from a few highly-rewarding regions, or to visiting as many different regions as possible.

The science hypothesis map integrated with this formulation in equation 7 can accommodate many different path planning strategies. There exist straightforward methods that use mixed integer linear programming (MILP) to generate optimal informative paths, but only when the rewards between sampled points are independent [20]. There is a family of algorithms that use the submodular property of information gain to produce near-optimal solutions [21], [22]. There are also some Monte Carlo tree search (MCTS) methods for informative path planning [12]. Other approaches may be computationally intensive, but potentially closer to optimality, such as using branch and bound (B&B) techniques for both discrete [23] and continuous [24] space representations.

IV. GEOLOGIC EXPLORATION

Often planetary exploration, including surface exploration of Mars, the Moon, and other planetary bodies, involves mapping surface mineralogy to infer geologic composition, structure, origins, and ages. This paper considers a geologic exploration scenario in which:

- The investigation objectives H are *geologic units* defining distinct units of material with different ages and formation processes. Figure 4 shows an example of a hydrothermal formation.
- The physical properties Y are *minerals*, distinctive chemical compositions that diagnose the geologic formation conditions of rocks. Figure 4 shows that a hydrothermal formation is characterized by the presence of minerals such as chlorite and serpentine.
- The measurements Z , *reflectance spectra*, representing the fraction of incident light reflected in each wavelength from the visible to shortwave infrared. Many minerals' molecular compounds have distinctive features in this range due to their distinctive chemical structure [25], as may be observed in Figure 4.

In order to train and test the science hypothesis map in a geologic mapping scenario, this paper focuses on the Cuprite mining district of Nevada, a region with rich mineral




Fig. 4: Visual example of the relation between geologic units, minerals and reflectance spectrum measurements. Elaborated with information from the United States Geological Survey (USGS) [26].

content and extensive field study [27]. This section describes the followed procedure to train the hierarchical probabilistic model, i.e. learn the association between geologic units and minerals, $P(Y|H)$, and the relation between minerals and reflectance spectrum measurements, $P(Z|Y)$.

A. Relation Between Geologic Units and Minerals

Figure 5 shows maps corresponding to the different elements of the science hypothesis map. Geologic units H and minerals Y are based on expert-drawn geologic and mineral maps by Swayze et al [27], which make a rough segmentation of the whole Cuprite site into regions. These maps are adjusted with manual control points so that they align on a per-pixel basis. Each independent geologic unit label and mineral are assigned a different value, illustrated as arbitrary false colors. These maps provide a ground-truth interpretation for each pixel in the scene. Note that each unit generally contains many minerals in different proportions. These proportions are used to train a conditional probability table of minerals given geologic classes, i.e. $P(Y|H)$. This table consists of 32 geologic units and 20 minerals. The training set uses up to 2000 random samples from every predefined region. These datasets are withheld from the analysis that follows.

B. Relation Between Minerals and Reflectance Spectra

Figure 5 also shows representative reflectance spectra of some key minerals in Cuprite. In-situ measurements performed by a robot with a spectrometer are simulated using an airborne instrument. Specifically, data from the Next Generation Airborne Visible Infrared Imaging Spectrometer (AVIRIS-NG) [28]. It assigns a unique reflectance spectrum measurement to every location (pixel) in the scene. AVIRIS-NG mapped the area at high spatial resolution (3.9 m per pixel) with radiance measurements from 380-2510 nm, also with a high spectral resolution (5.0 nm per channel). The data was acquired during overflights in 2014 and converted from measured at-sensor radiance to surface reflectance using the procedure described by [29].




Fig. 5: Cuprite, NV. Left: geologic units H . Center: minerals Y and their abundances. Right: Representative reflectance spectra Z of some key minerals. Each pixel is associated with a full spectrum. Geologic units and mineral abundances are based on the work by Swayze et al [27].

A Gaussian Mixture Model (GMM) is used as a simple but effective probabilistic classifier that predicts the corresponding mineral from a spectrum measurement. The model is trained using a held-out training set of locations from each mineral class. Since the used AVIRIS dataset has a very high resolution, its dimensionality is reduced to 20 dimensions with Principal Component Analysis (PCA). Then, a mean and covariance matrix is fit to each class using the shrinkage estimator described in [30]. The shrinkage estimator treats the actual covariance as a linear combination of the sample covariance, and a regularized version consisting of the diagonal elements only. A closed form solution to the leave-one-out cross validation (LOOCV) likelihood allows efficient computation of the optimal linear combination. The end result is a probability density for every combination of mineral and spectrum in the image.

Treating Z as a discrete variable simplifies information gain calculations. Consequently, the classification output of the GMM classifier is treated as a categorical answer. Figure 6 shows the corresponding confusion matrix from a held-out validation set to infer its probability of error on any future measurement. It has an overall accuracy of 57.15%, with a min/max class accuracy of 12.50% and 98.95%, respectively. Although it is a classifier with a relatively low accuracy, its confusion matrix acts as a conditional probability table that compensates the noise in the measurement process which could cause the estimated mineral Z to differ from the true physical variable Y . The used training and validation sets each use up to 1000 samples from every mineral type. These datasets are withheld from the analysis that follows.

V. EXPERIMENTS: BAYES LEARNING

This group of experiments evaluated the model's ability to predict and recover the true geologic unit of a region (with a specific age, formation process, etc.) from high-resolution AVIRIS spectral measurements with the Bayesian update. For each of the 32 regions, 200 random sampling sequences were generated and averaged. Non-training points were sampled without replacement, simulating an exhaustive exploration with no budget constraints.




Fig. 6: Empirical conditional probability table relating actual to estimated minerals for each of 20 possible mineral types in the Cuprite scene.

To evaluate the model's ability to recover the true geologic unit of each region, this experiment tracked two fundamental variables: the posterior probability of the correct geologic unit, and the updated entropy. As more and more measurements are collected, the updated probability of the true underlying geologic unit should ideally converge to 1, and consequently, the entropy should converge to 0. These convergence rates depend on the prior distribution, therefore, different initial conditions were tested through three representative scenarios (Figure 7), explained next.

Starting with an accurate prior: To represent a situation where the scientist has a relatively accurate initial belief, a value of 50% was assigned to the prior probability of the correct geologic unit, while the remaining probability mass is distributed to alternatives. This prior allows ample margin for converging to either a correct or incorrect answer. Figure 7 shows that most beliefs converge to the right answer with few samples, while entropies reduce significantly. The three exceptions with poor outcomes correspond to the smallest regions trained with the fewest data points.

Starting with complete uncertainty: In this case, a situation of complete uncertainty was represented with a uniform prior over geologic unit labels. Despite the more challenging situation, Figure 7 shows that most beliefs still show significant




Fig. 7: Posterior probability of the correct class (top row) and posterior entropy (bottom row), starting with an accurate prior (left column, green medians), under complete uncertainty (center column, blue medians), and an inaccurate prior (right column, red medians). Each data point in the box and whisker plots corresponds to one of the 32 different regions.

improvement. As expected, there is more variance and an overall slower convergence rate.

Starting with an inaccurate prior: Finally, this part represents a situation where the scientist’s initial beliefs are actually incorrect. For that, a prior probability of 50% was assigned to an incorrect unit: the most similar geologic unit according to the Hellinger distance for probability distributions, a challenging error to correct. The remaining probability mass is distributed uniformly. Figure 7 shows an even slower improvement with a higher variance. But in most cases the model still recovers.

This results indicate that the model is able to recover the true geologic unit most of the times as spectral measurements are collected, even under misleading initial conditions.

VI. EXPERIMENTS: PATH PLANNING

This group of experiments integrated the science hypothesis map with different information gain-driven path planners. It compares their performance in order to evaluate the potential benefits in science yield when using information gain-based planning over conventional science-blind planners.

A. Experimental Setup

The whole Cuprite region was divided into three sites with the same area: A, B, and C. For simplicity, each site is treated as a 2D discrete graph that is formed by an 8-connected rectangular grid, with locations spaced at 50 pixels (200 m). The full map has a size of 2555×2268 pixels (10.22×9.07 km). For each site, 200 different pairs of random start and end points were generated, located at the East and West edges of the map. The cost function was directly proportional to




Fig. 8: Example paths from the 6 path planners. The start is in the upper left, and the end in the lower right.

the path length, and the assigned budget permitted up to 1.5 times the shortest path length, allowing sufficient budget to explore alternative regions. All simulated traverses and regions had the same initial conditions: a prior probability of 40% assigned to both the correct geologic unit and the most similar unit according to the Hellinger distance, while the remaining probability mass was distributed to alternatives. This represented a challenging scenario where the scientist is equally inclined toward two similar options.

B. Tested Path Planners

Figure 8 shows an example of paths generated by all the tested planners. As a control case, this paper evaluates three science-blind planners that ignore information gain:

- 1) A *Direct* (D) path planner that selects the sequence of waypoints that minimizes the total path length with the standard Dijkstra algorithm.
- 2) A *Random* (R) path planner that sequentially selects a set of random neighboring waypoints that do not violate the budget on path length.
- 3) A *Max-N* (M) path planner that maximizes the number of visited nodes using a classic recursive greedy approach for orienteering [31]. Since there may be multiple valid solutions, it performs a random permutation of the waypoints’ ids in order to avoid biased paths (e.g. paths that prioritize going to the left).

Against these, this paper evaluates three science-aware planners that incorporate measurements’ information gain:

- 1) A *Greedy* (G) path planner that adds at each step the neighboring waypoint that maximizes the objective function, this without exceeding the budget.
- 2) A *Branch and bound* (B) path planner as a nonmyopic alternative. It selects a sequence of waypoints using the algorithm described by Binney and Sukhatme [23]. In this particular case, implementing a three-step calculation look ahead.
- 3) A *pSPIEL* (P) planner that combines submodular orienteering algorithms by Singh et al. [22] with this paper’s

TABLE II: Average scores for three relevant path planning metrics: 1) path length in meters, 2) number of collected measurements, 3) number of explored regions.

Site	Metric	D	R	M	G	B	P
A	1	3937	5582	6129	5639	5815	5817
	2	15.99	23.24	27.67	23.43	24.94	26.74
	3	5.32	5.95	6.03	7.71	8.20	8.77
B	1	3789	5430	5892	5368	5511	5608
	2	13.34	19.88	23.66	18.82	20.44	22.51
	3	5.62	6.32	6.38	8.05	8.83	9.65
C	1	4045	5798	6285	5468	5937	5974
	2	12.29	18.77	21.67	17.49	20.31	21.32
	3	3.94	4.38	4.49	5.24	5.53	5.88

information gain objective function.

Additionally, these planners updated the hypotheses and paths every time a new measurement was collected, this with the goal to simulate an explorer that adapts its path dynamically with each new observation.

C. Results

Table II shows the average scores per traverse of three simple but useful metrics: path length, number of collected measurements, and number of explored regions. As expected, the Direct planner minimizes the first two metrics, whereas the Max-N planner maximizes them. The rest of the planners have an intermediate performance, with the Random and Greedy planners getting similar scores, and the pSPIEL planner achieving the second best performance overall. Nonetheless, all the science-aware planners explore more regions than any of the science-blind planners, apparently favoring exploration over exploitation in this scenario.

The performance of the planners is measured from an information-theoretic perspective with the following variables: the evolution of the posterior probability and its entropy. The evaluation of a path across the map is done by simply adding the corresponding updated region’s metrics, where a poor performance in a region penalizes the global score, and vice versa. These scores have a high variance because they strongly depend on the assigned path length budget. Therefore, the scores are normalized using feature scaling with respect to the best and worst planners for each given pair of start-end locations. The resulting scaled scores for prediction accuracy and entropy are shown in Figure 9. On the other hand, Table III shows a set of paired t-tests comparing the performance among planners for both metrics with their raw unnormalized scores.

These box plots and hypothesis tests show a couple of things. The science-blind planners have a performance more or less proportional to the number of measurements, being Max-N the best. However, that is not necessarily true when these are compared to the science-aware planners: they tend to achieve superior scores in both of the information-theoretic metrics, even when some of them spend less budget or collect less measurements in average (e.g. Greedy vs. Random and Max-N).




Fig. 9: Comparison of planners based on their ability to recover correct geologic units (top row) and the final entropy (bottom row) per traverse using a normalized score. The green dots are the means.

This shows that there are more meaningful science measurements than others, and balancing them adequately leads to higher science productivity. For instance, the nonmyopic planners (B&B and pSPIEL) outperform the Greedy method since it only has a one-step look ahead. pSPIEL is the best overall method because it analyzes and weights all reachable regions using an approximation graph, favoring the ones with the highest expected rewards. On the other hand, the B&B method is limited by a three-step calculation horizon.

VII. CONCLUSION

Recent robotics advances in the areas of long-duration autonomy and long-distance mobility illustrate the need for this paradigm shift in human co-robot exploration. Robots can now spend days in the field and travel multiple kilometers per single uplink/downlink communications cycle. These robots can deploy instruments over wide areas, augmenting human explorers and acting on their behalf when appropriate. Our future work will continue to investigate this problem as a venue to refine the relationship between human scientists and robotic explorers, ultimately increasing the productivity and yield of both.

We highlight several key findings of this initial investigation. First, the experiments demonstrate that the science hypothesis map is a sound mathematical framework for describing exploration for specific well-defined objectives. The probabilistic model can infer geologic units with a high accuracy under diverse challenging situations. In the case of robotic path planning, path planners that exploit information gain, efficiently calculated with the GLF approximation, better reduce uncertainty over the investigation objectives. Not surprisingly, nonmyopic planning outperforms greedy and science-blind alternatives.

TABLE III: Set of paired t-tests comparing the performance among planners using a significance level of 5%. In most cases, there is a statistically significant difference (Y/N) in both inference (first entries) and entropy (last entries).

Site Planner	A						B						C					
	D	R	M	G	B	P	D	R	M	G	B	P	D	R	M	G	B	P
D	-	Y, Y	Y, Y	Y, Y	Y, Y	Y, Y	-	N, Y	N, Y	Y, Y	Y, Y	Y, Y	-	Y, Y	Y, Y	Y, Y	Y, Y	Y, Y
R	-	-	Y, Y	Y, Y	Y, Y	Y, Y	-	-	N, N	Y, Y	Y, Y	Y, Y	-	-	N, Y	Y, Y	Y, Y	Y, Y
M	-	-	-	Y, Y	Y, Y	Y, Y	-	-	-	Y, Y	Y, Y	Y, Y	-	-	-	N, N	Y, Y	Y, Y
G	-	-	-	-	Y, Y	Y, Y	-	-	-	-	N, Y	Y, Y	-	-	-	-	Y, Y	Y, Y
B	-	-	-	-	-	N, Y	-	-	-	-	-	Y, Y	-	-	-	-	-	Y, Y
P	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

A more surprising finding is the strength of the statistical link between measured spectra, the corresponding mineral types, and the geologic unit classes. This is not a foregone conclusion because geologists form this classification from many other features such as local morphology, elevation, the wide-area geographic three-dimensional structure of different strata, and domain knowledge. It is striking that AVIRIS-NG reflectance spectra, drawn independently and randomly from each unit, so strongly predict the units' unique type out of over thirty alternatives. Future work will improve the fidelity and realism of both the constraints and the measurement data, while incorporating scientist revisions on multiple sequential command cycles.

VIII. ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation National Robotics Initiative Grant #IIS-1526667. We acknowledge support of the NASA Earth Science Division for the use of AVIRIS-NG data. A portion of this work was performed at the Jet Propulsion Laboratory, California Institute of Technology. Copyright 2017, All Rights Reserved. U.S. government support acknowledged.

REFERENCES

- [1] A. N. Hock *et al.*, "Life in the atacama: A scoring system for habitability and the robotic exploration for life," *J. Geophysical Research: Biogeosciences*, vol. 112, no. G4, 2007.
- [2] T. A. Estlin *et al.*, "Aegis automated science targeting for the mer opportunity rover," *ACM TIST*, vol. 3, no. 3, p. 50, 2012.
- [3] D. R. Thompson *et al.*, "Autonomous science during large-scale robotic survey," *J. Field Robotics*, vol. 28, no. 4, pp. 542–564, 2011.
- [4] J. Binney *et al.*, "Informative path planning for an autonomous underwater vehicle," in *Proc. IEEE ICRA*, 2010, pp. 4791–4796.
- [5] K. Chaloner and I. Verdine, "Bayesian experimental design: a review," *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995.
- [6] M. Woods *et al.*, "Autonomous science for an exomars roverlike mission," *J. Field Robotics*, vol. 19, no. 5, pp. 467–482, 2009.
- [7] S. Chien *et al.*, "Using autonomy flight software to improve science return on earth observing one," *J. Aerospace Computing, Information, and Communication*, vol. 2, pp. 196–216, 2005.
- [8] A. Castano *et al.*, "Automatic detection of dust devils and clouds on mars," *Machine Vision and Applications*, vol. 19, no. 5, pp. 467–482, 2008.
- [9] M. Gallant *et al.*, "Science-influenced mobile robot guidance using bayesian networks," in *Proc. Canadian Conf. on Electrical and Computer Engineering*, 2011, pp. 1135–1139.
- [10] L. Pedersen *et al.*, "Autonomous robotic meteorite identification in antarctica," in *Proc. IEEE ICRA*, 2001, pp. 4158–4165.
- [11] Y. Girdhar and G. Dudek, "Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring," *Autonomous Robots*, vol. 40, no. 7, pp. 1267–1278, 2016.
- [12] A. Arora *et al.*, "Extending autonomy of planetary rovers by encoding geological knowledge in a bayesian framework," in *10th International Cognitive Robotics Workshop*, 2016.

- [13] D. Lindley, "On the measure of information provided by an experiment," *The Ann. of Math. Stat.*, vol. 27, no. 4, pp. 986–1005, 1956.
- [14] J. Bernardo, "Expected information as expected utility," *The Ann. of Stat.*, vol. 7, no. 3, pp. 686–690, 1979.
- [15] K. Ryan, "Estimating expected information gains for experimental designs with application to the random fatigue-limit model," *J. of Computational and Graphical Statistics*, vol. 12, no. 3, pp. 585–603, 2003.
- [16] J. van den Berg *et al.*, "Optimal nonlinear bayesian experimental design: an application to amplitude versus offset experiments," *Geophysical J. International*, vol. 155, no. 2, pp. 411–421, 2003.
- [17] D. Haussler and M. Opper, "General bounds on the mutual information between a parameter and n conditionally independent observations," in *Proc. 8th Annu. Conf. Computational Learning Theory*, 1995, pp. 402–411.
- [18] F. Richards, "A flexible growth function for empirical use," *J. Experimental Botany*, vol. 10, no. 2, pp. 290–301, 1959.
- [19] N. A. Smith and R. W. Tromble, "Sampling uniformly from the unit simplex," Johns Hopkins Univ., Baltimore, MD, Tech. Rep. 29, 2004.
- [20] N. Yilmaz *et al.*, "Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming," *IEEE J. Ocean. Eng.*, vol. 33, no. 4, pp. 522–537, 2008.
- [21] A. Krause *et al.*, "Near-optimal sensor placements: maximizing information while minimizing communication costs," in *In ACM/IEEE IPSN*, 2006, pp. 2–10.
- [22] A. Singh *et al.*, "Nonmyopic adaptive informative path planning for multiple robots," in *Proc. of IJCAI*, vol. 34, 2009, pp. 1843–1850.
- [23] J. Binney and G. S. Sukhatme, "Branch and bound for informative path planning," in *Proc. of IEEE ICRA*, 2012, pp. 2147–2154.
- [24] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The Int. J. of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [25] R. N. Clark, "Spectroscopy of rocks and minerals, and principles of spectroscopy," *Manual of Remote Sensing*, vol. 3, pp. 3–58, 1999.
- [26] R. Clark *et al.*, "Usgs digital spectral library splib06a," *U.S. Geological Survey, Digital Data Series 231*, 2007.
- [27] G. A. Swayze *et al.*, "Mapping advanced argillic alteration at cuprite, nevada, using imaging spectroscopy," *Economic Geology*, vol. 109, no. 5, pp. 1179–1221, 2014.
- [28] L. Hamlin *et al.*, "Imaging spectrometer science measurements for terrestrial ecology: Aviris and new developments," in *Aerospace Conference, 2011 IEEE*. IEEE, 2011, pp. 1–7.
- [29] D. R. Thompson *et al.*, "Atmospheric correction for global mapping spectroscopy: Atrem advances for the hyspirci preparatory campaign," *Remote Sensing of Environment*, vol. 167, pp. 64–77, 2015.
- [30] J. P. Hoffbeck and D. A. Landgrebe, "Covariance matrix estimation and classification with limited training data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 7, pp. 763–767, 1996.
- [31] C. Chekuri and M. Pal, "A recursive greedy algorithm for walks in directed graphs," in *Proc. IEEE FOCS*, 2005, pp. 245–253.